

Use of MPI_File_set_view

MPI_File_set view is a collective call, and we don't like having to call it as often as we do. However, that's the way MPI is written, so we end up needing nearly every MPI_File_read_all and MPI_File_write_all call to partner with an MPI_File_set_view.

We also reset the file view after calling an MPI I/O call: we need to do this so that subsequent independent i/o calls know the state of the file view.

Why can't we reset the file view in ncmpi_begin_indep_data ? Because MPI_File_set_view is collective.

We do try to limit the number of MPI_File_set_view calls: we can avoid setting the fileview entirely when the request is a single contiguous file region. In this case, MPI-IO calls with explicit offset functions are used.